

**IN THE UNITED STATES DISTRICT COURT  
EASTERN DISTRICT OF VIRGINIA  
ALEXANDRIA DIVISION**

AMDOCS (ISRAEL) LIMITED, an Israeli  
Corporation,

Plaintiff,

v.

OPENET TELECOM, INC., a Delaware Corporation,  
and OPENET TELECOM LTD., an Irish Corporation,

Defendants.

Case No. 1:10cv910 (LMB/TRJ)

**DECLARATION OF JOSEPH HOGAN**

I, Joseph Hogan, hereby declare as follows:

1. I am the co-founder and Chief Technology Officer (“CTO”) of Defendant Openet Telecom Ltd. I am submitting this declaration in connection with Openet’s Motion for Summary Judgment.

2. Openet Telecom Ltd. was spun out of Openet International in 1999, which was founded in 1995. As co-founder and CTO of Openet Telecom Ltd., I have been personally involved in most key aspects of the design and development of Openet’s products, including the FusionWorks products at issue in this case.

3. I have first-hand knowledge of the matters set forth in this declaration. Further, I was deposed for two days by Amdocs in this case, and if this case proceeds to trial, I anticipate testifying at trial.

**Overview of the Accused Infringing FusionWorks Products**

4. I understand that Amdocs accuses FusionWorks Mediation and products sold in conjunction with the FusionWorks Framework, including Convergent Mediation, Convergent

Charging, Network-Edge Rating, Balance Manager, Profile Manager, and Policy Manager, of infringing the asserted claims of the '065, '797, and '510 patents. I understand that only FusionWorks Mediation and Convergent Mediation are accused of infringing the asserted claims of the '984 patent.

5. Given these allegations, it is necessary to clarify how the Framework relates to the individual products offered by Openet. To that end, I have attached as Exhibit 1 to this declaration a copy of select excerpts of FusionWorks Convergent Mediation User Guide (Software Release 6.2), dated October 21, 2010 (hereinafter, "the User Guide").

6. The Framework is not a standalone product. The Framework itself does not perform any functions and is inert without additional software and configuration. Rather, the Framework is a collection of basic components. Depending on how the product is configured, different products may use some of these components. These components may include tools for system monitoring, configuration, and security. *See* OPENENT 161581.


7. Each component within the Framework operates according to business logic rules written for that component. Business logic rules are written using DataStream Decoder (DSD) scripts.

8. One component of the Framework is the Correlation and Transaction Engine ("CTE"). The CTE is a rules-driven engine. In other words, the CTE only operates according to business logic rules (or DSD scripts) that have been written for the CTE. Such rules vary from customer to customer and from installation to installation.

9. FusionWorks Convergent Mediation (and its predecessor product, FusionWorks Mediation) is a software product that mediates (*i.e.*, analyzes and processes) data generated by a telecommunications network after a service has been delivered (*e.g.*, after a call has been made, a

text message has been sent, or a webpage has been downloaded). The mediated data is then passed downstream to a billing system.

10. Convergent Mediation includes two main components, the Mediation Server and Collectors, which operate in connection with the Framework, as depicted in the following excerpt from the User Guide.



11. Thus, while mediation functions are performed within the Framework, the Framework must be configured with the necessary components and with necessary DSD code to perform mediation. In other words, the Framework does not natively perform mediation.

12. The other accused products (Convergent Charging, Network-Edge Rating, Balance Manager, Profile Manager, and Policy Manager) are plug-ins to the Framework. Thus, unlike Convergent Mediation, the functionalities of Convergent Charging, Network-Edge Rating, Balance Manager, Profile Manager, and Policy Manager are performed outside the Framework.

13. Convergent Charging, Network-Edge Rating, Balance Manager, Profile Manager, and Policy Manager use certain components of the Framework, such as the security or

configuration components, but the Framework does not perform mediation when operating with those products, nor is the Framework deployed with the DSD code needed to perform mediation when it is supporting those products.

14. More fundamentally, Convergent Charging, Network-Edge Rating, Balance Manager, Profile Manager, and Policy Manager are not mediation systems. While Convergent Mediation processes data generated by a telecommunications network once a service has been rendered, Convergent Charging, Network-Edge Rating, Balance Manager, Profile Manager, and Policy Manager are considered “transactional intelligence” products that allow a provider of telecommunications services to manage access to services as they are delivered. For example, Policy Manager allows a wireless carrier to check whether a subscriber is entitled to access certain content or services, Convergent Charging allows carriers to apply charges to customers when exceeding monthly usage limits, and Network Edge Rating allows carriers to adjust rates charged for services based on demand and usage. Transactional intelligence is a relatively new field that did not exist at the time of the filing patents-in-suit.

#### **Alleged Infringement of ‘797 Patent**

15. U.S. Patent No. 6,836,797 (“the ‘797 patent”) claims various steps, including “collecting data describing the plurality of services” and “generating a single record represent[ing] each of the plurality of services.” As an initial matter, based on my experience in the telecommunication software industry, some people in the late 1990s may have perceived generating a single record representing each of the plurality of services as a processing efficiency (although I believe that generating a single record would not significantly improve efficiencies in the overall process of mediation and billing). However, increases in computing power and data storage quickly eliminated any incentive to consolidate data into a single record.


16. For the reasons explained below, no FusionWorks product generates a single record representing each of the plurality of services. Doing so would require writing DSD code to generate such a record, and I am aware of no such code.

17. FusionWorks Convergent Mediation gathers information from a network and converts the information into a generic data structure called an Attribute Value Stream (“AVS”). AVS data fields can be grouped together to create an “event.”

18. Fusion Works Convergent Mediation includes an “Event Editor” that allows a user to define the data fields to be included in some events. The fields must be defined before data is collected.

19. FusionWorks Convergent Mediation includes a graphic user interface (“GUI”) that allows the user to select “rules” that are applied to an event when the event is processed, as depicted below in the following excerpt from the User Guide:

*Figure 8-5 Enter Mapping Dialog*



20. Rules are applied to all fields in an event. The GUI does not allow rules to be applied to specific fields within an event. FusionWorks lacks the claimed requirements of listing functions, choosing at least one field, and applying at least one listed function to the chosen field

21. When events are processed, events of the same service type (e.g., voice, data) can be consolidated into one record. Records reflecting different services can be aggregated into a file containing multiple records of the different service types but cannot be consolidated into a

single record representing each of the services unless DSD code is written to do so, but no such code has been written.

### **Alleged Infringement of ‘984 and ‘510 Patents**

22. U.S. Patent Nos. 6,947,984 (“the ‘984 patent”) and 7,412,510 (“the ‘510 patent”) claim methods and computer code for “reporting on a collection of network usage information.” As an initial matter, similar to the ‘797 patent, the claimed technology of the ‘984 and ‘510 patents is inapposite to FusionWorks Convergent Mediation. The claims of both patents require “storing [a] plurality of data records in a database” and “submitting queries to the database using the [selected or predetermined] reports to for retrieving information on the collection of network usage information from the network devices.” Such an application may have some utility in the billing context where records are stored until a bill is generated but is typically not performed by a mediation system. Indeed, FusionWorks Convergent Mediation mediates records and passes them downstream to the billing system and does not store data records in a database for later retrieval.

23. Amdocs apparently accuses the statistics system in FusionWorks of infringing these patents. The “statistics system . . . enables statistics definition, delivery, and storage.” *See* User Guide at OPENET 164777.

24. The statistics system does not enable reporting on the collection of network usage information at each network device; instead, the statistics system reports on the internal operation of the FusionWorks processors (*e.g.*, “number of CDRs processed within a period” by FusionWorks, *see* OPENET 164788). Statistics would, for example, represent the number of records collected from a switch during a 24-hour period, the number of bad records collected during a period, or the number of records that failed validation. This information may be used to

diagnose the operation of FusionWorks Convergent Mediation but is not useful for reporting or billing on network usage.

25. Further, the statistics system does not natively collect statistics and generate reports; instead additional computer code – written in DataStream Decoder (DSD) format – is required to collect statistics and generate reports. *See* OPENET 164777 (“[A] DSD programmer can collect and send statistics to the statistics system.”).

26. I am not aware of such DSD scripts written by Openet for any U.S. customer. In my experience, U.S. customers prefer to use their own reporting systems rather than the system provided by Openet within FusionWorks. To the extent U.S. customers use the Openet statistics system, they write their own DSD scripts.

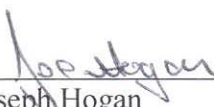
27. Additionally, FusionWorks does not process events according to a fixed latency. *See* OPENET 164789 (“average latency” and “minimum and maximum latency for a particular section of event processing in a period” vary over time).

28. Because system latency varies, FusionWorks does not ensure a fixed latency. Because latency varies, at any given time the collection of events can exceed the latency of the system.

[rest of page left intentionally blank]

I declare, subject to the penalty of perjury of the laws of the United States of America, that the foregoing is, to the best of my knowledge, true and correct.

Date: May 26, 2011  
Dublin, Ireland

  
\_\_\_\_\_  
Joseph Hogan



# Exhibit 1



## **FusionWorks Convergent Mediation User Guide**

Software Release 6.2  
October 21, 2010

### **Corporate Headquarters**

Openet  
6 Beckett Way  
Park West Business Park  
Dublin 12  
Ireland  
<http://www.openet.com>  
Tel: +353 1 620 4600  
Fax: +353 1 620 4990



Intelligence

# CHAPTER 1

## System Overview

---

This chapter provides an overview of FusionWorks Convergent Mediation and the FusionWorks Framework.

It outlines the main components of FusionWorks Convergent Mediation, including the Correlation and Transaction Engine (CTE) and the collector. It also provides a detailed description of feature changes and new functionality in this release of FusionWorks Convergent Mediation.

It includes the following sections:

- FusionWorks Framework Overview, page 1-1
- FusionWorks Framework Components, page 1-2
- FusionWorks Convergent Mediation Components, page 1-5
- Role of FusionWorks Convergent Mediation, page 1-8

## FusionWorks Framework Overview

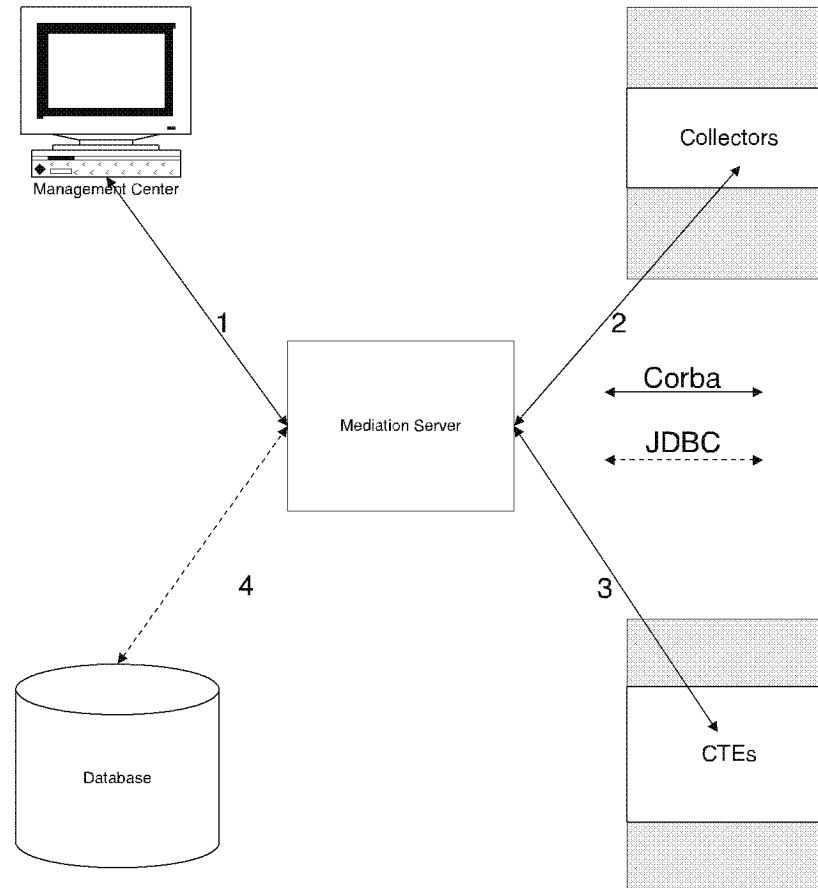
The FusionWorks Framework provides a common foundation for the Openet suite of convergent network-edge products.

FusionWorks Framework components allow for the re-use of common functionality such as database management, configuration, statistics reporting and security/user management.

See Figure 1-1 that outlines the FusionWorks Framework.

**Figure 1-1**      **FusionWorks Framework**



**Figure 1-3 Mediation Server Process Interaction**

## Component Independence

Component independence is a feature which simplifies the configuration of FusionWorks Framework in a highly available architecture. It allows traffic processing components (e.g., collectors and CTEs) to continue operating in a largely unimpaired capacity, even if certain supporting components of the system are unavailable. This behavior means that clustering system restart rules are simplified. For example, CTEs and collectors no longer need to be restarted following a Mediation Server restart. This feature also simplifies the general operation of the system.

## Supporting Components

Supporting components include the Mediation Server, the Business Rule Manager, the alarm manager, unified logging, and the statistics server.

## AVS Connections for CTEs

A CTE can have multiple AVS protocol sessions open at any time, but it cannot have more than one outward session open with any one particular component at any time. It can have multiple inward sessions, but only one outward session which can be written to the permanent store.

## Mapping Events to Rules

There are two separate mapping options available to the user:

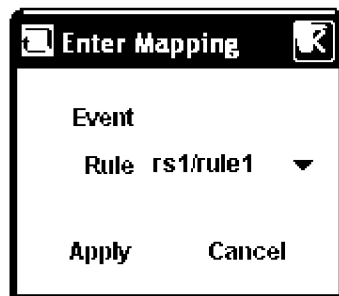
- Mapping Incoming AVS Events to Rules, page 8-9
- Mapping Other Events to Rules, page 8-9

### Mapping Incoming AVS Events to Rules

Use the following steps to map an incoming AVS event to a rule (DSD script):

- Step 1** Double-click the Incoming AVS Event you want to associate with a rule. A dialog appears. (See Figure 8-5.)

**Figure 8-5** Enter Mapping Dialog



- Step 2** Choose a rule from the list of available rules in the drop-down list.
- Step 3** Click **Apply** to commit the changes or **Cancel** to exit.  
The new rule is displayed in the list of rules in the Incoming AVS Events list, paired with the event to which it is mapped.  
The rule will be run on the CTE when the particular AVS event is received.



**Note**

A CTE will not start if any of the rules associated with its events are missing. This might occur after an upgrade for example. You can examine the logs to determine if any rules are missing from the workspace. You must restore the missing rules and restart the CTE.

### Mapping Other Events to Rules

To map a system event or UDE to a rule (DSD script), follow these steps:



# CHAPTER 15

## Statistics

---

This chapter details the statistics system, which enables statistics definition, delivery and storage. This chapter explains how to define statistics and statistic key performance indicators (KPIs) using the Management Center, and describes how a DSD programmer can collect and send statistics to the statistics system. It also gives an overview of viewing statistics and KPIs in the OAM Center.

Statistics can also be delivered over SNMP. This functionality is not detailed in this chapter, but is described in Chapter 18, “SNMP Agent Management”. See the “Delivering Statistics over SNMP” section on page 18-9 for more information.

**Note**

The functionality described in this chapter relates only to the new statistics system which was introduced in FusionWorks Framework 5.0, that is, it does not describe the statistics and reporting facilities that existed in previous releases. However, those facilities are still available in this release and are described in Appendix A, “Statistic Reports”.

---

The sections contained in this chapter are as follows:

- Overview, page 15-1
- Intermediate File Storage, page 15-8
- Statistics Categories, page 15-11
- Statistic Definitions, page 15-14
- Configuring the Statistics System, page 15-38
- Configuring the Statistics System in the OAM, page 15-45
- Viewing Statistic KPIs in the OAM, page 15-49
- Statistics DSD API, page 15-55

**Note**

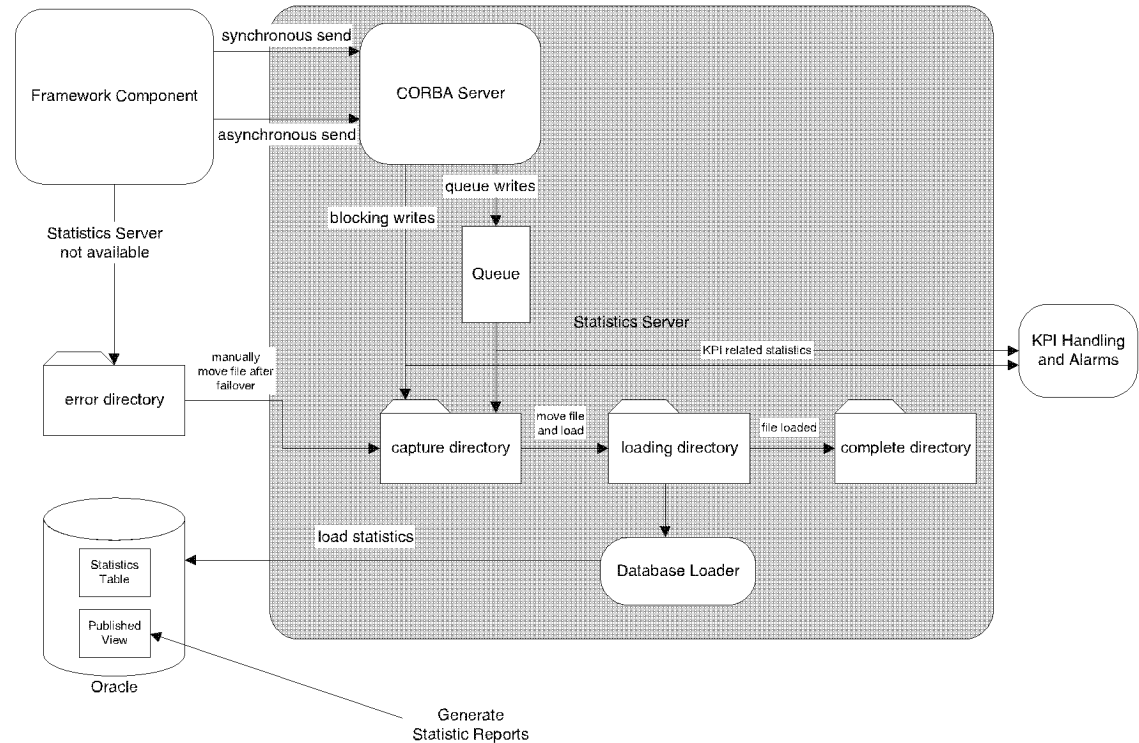
You can optionally use a separate database schema store for statistics. See *FusionWorks Framework Installation Guide* for more information.

---

## Overview

This section provides an overview of the statistics system. It contains the following subsections:

- Terminology, page 15-2
- Features and Benefits, page 15-3

**Figure 15-4 IFS File Lifecycle for Statistics**

## Statistics Categories

This implementation of statistics defines several statistics measurement categories, which describe how the statistics are processed at runtime. When creating a statistic definition, one of these statistics categories must be specified. The supported categories are described in the following sections:

- Periodic Count, page 15-11
- Simple Count, page 15-12
- Associated Count, page 15-12
- Periodic Performance Measure, page 15-13
- Latency Measure, page 15-13
- Periodic Sample, page 15-13
- Single Event, page 15-14

There is also the concept of statistics groups, which is explained in the “Groups” section on page 15-14.

## Periodic Count

This is a numeric count associated with a period. This is a single valued statistic.

## Statistic Value Type

A periodic count supports numeric value types only. Values start at 0.

## Examples

The following are examples of periodic counts:

- Number of CDRs processed within a period.
- Number of CDRs of a particular type within a period.
- Number of protocol errors in a period.

## Simple Count

This is very similar to the periodic count category, except it does not reset the value when the configured period expires. This is a single valued statistic.

## Statistic Value Type

A simple count supports numeric value types only. Values start at 0. The maximum value is 2147483647. If the maximum value is exceeded the behaviour is undefined.

## Examples

The following are examples of simple counts:

- Count of successful transactions.
- Count of poorly formed requests received.

## Associated Count

This is a numeric count associated with a named object (for example, a file). This is a single valued statistic.

## Statistic Value Type

An associated count supports numeric value types only. Values start at 0.

## Examples

The following are examples of associated counts:

- Number of CDRs in a batch file.
- Number of CDRs of a particular type in a batch file.
- Number of duplicates found in a batch file.



## Periodic Performance Measure

This is a transactional performance measure associated with a period. This is multi-valued statistic containing average latency, maximum latency, minimum latency, transactions per second, transaction ID of the maximum latency event, or total transactions.

### Statistic Value Type

A periodic performance measure supports numeric and string value types (not user specifiable).

### Examples

The following are examples of periodic performance measures:

- Average latency for a particular section of event processing in a period.
- Minimum and maximum latency for a particular section of event processing in a period.
- The transactions per second for a particular section of event processing in a period.
- The IDs of the transactions that had the maximum latencies within a period.
- Number of transactions exceeding a latency threshold in a period as defined by SLA.

## Latency Measure

This is a series of latency measures, each of which may have a transaction ID. This is a single valued statistic.

### Statistic Value Type

A latency measure supports numeric value types only.

### Examples

The following are examples of latency measures:

- Time it took to collect (FTP) a file.
- Time it took to process a single file.
- Time it took to process a batch of files.

## Periodic Sample

This is a periodic sample of the instantaneous value of some state value. This is a single valued statistic.

### Statistic Value Type

A periodic sample supports numeric or string value types only.

## Examples

The following are examples of periodic samples:

- The current size of a component's threadpool queue.
- The name of the batch file currently being processed.

## Single Event

This is an aperiodic event containing any arbitrarily generated value of interest. It may or may not be associated with an object. This is a single valued statistic.

## Statistic Value Type

A single event supports numeric or string value types only.

## Examples

The following are examples of single events:

- The time at which a particular batch file was processed.
- The checksum of an input file.
- The first and last CDR sequence number in a file (two separate statistic instances required).
- The list of billing output files that were generated from a collected CDR file (one statistic instance per billing output file).

## Groups

Groups are used to maintain a grouping of associated count type statistics whose counts can be checkpointed, committed and rolled back together. They are intended for use in batch environments where it is imperative that counts remain consistent across processing interruptions and resumptions mid-batch.

## Statistic Definitions

Statistic definitions are accessed via the Statistics Management node in the Management Center. There are four types of statistic definitions, namely:

- Custom Statistics—These are customised statistics which are highly specific to a customer's business logic. You can create, edit, and delete custom statistics.
- Product Statistics—These are the core statistics which are applicable to all deployments. Definitions for these statistics are created at installation time and they cannot be created or deleted. You can change the period (frequency at which they are collected), collection (if they are collected or not), persistence (if they are persisted or not), publication (if they are published or not), and lifetime (how long they are stored in the database) of the existing product statistic definitions.

## Statistic Value Type

A periodic count supports numeric value types only. Values start at 0.

## Examples

The following are examples of periodic counts:

- Number of CDRs processed within a period.
- Number of CDRs of a particular type within a period.
- Number of protocol errors in a period.

## Simple Count

This is very similar to the periodic count category, except it does not reset the value when the configured period expires. This is a single valued statistic.

## Statistic Value Type

A simple count supports numeric value types only. Values start at 0. The maximum value is 2147483647. If the maximum value is exceeded the behaviour is undefined.

## Examples

The following are examples of simple counts:

- Count of successful transactions.
- Count of poorly formed requests received.

## Associated Count

This is a numeric count associated with a named object (for example, a file). This is a single valued statistic.

## Statistic Value Type

An associated count supports numeric value types only. Values start at 0.

## Examples

The following are examples of associated counts:

- Number of CDRs in a batch file.
- Number of CDRs of a particular type in a batch file.
- Number of duplicates found in a batch file.

## Periodic Performance Measure

This is a transactional performance measure associated with a period. This is multi-valued statistic containing average latency, maximum latency, minimum latency, transactions per second, transaction ID of the maximum latency event, or total transactions.

### Statistic Value Type

A periodic performance measure supports numeric and string value types (not user specifiable).

### Examples

The following are examples of periodic performance measures:

- Average latency for a particular section of event processing in a period.
- Minimum and maximum latency for a particular section of event processing in a period.
- The transactions per second for a particular section of event processing in a period.
- The IDs of the transactions that had the maximum latencies within a period.
- Number of transactions exceeding a latency threshold in a period as defined by SLA.

## Latency Measure

This is a series of latency measures, each of which may have a transaction ID. This is a single valued statistic.

### Statistic Value Type

A latency measure supports numeric value types only.

### Examples

The following are examples of latency measures:

- Time it took to collect (FTP) a file.
- Time it took to process a single file.
- Time it took to process a batch of files.

## Periodic Sample

This is a periodic sample of the instantaneous value of some state value. This is a single valued statistic.

### Statistic Value Type

A periodic sample supports numeric or string value types only.

## Examples

The following are examples of periodic samples:

- The current size of a component's threadpool queue.
- The name of the batch file currently being processed.

## Single Event

This is an aperiodic event containing any arbitrarily generated value of interest. It may or may not be associated with an object. This is a single valued statistic.

## Statistic Value Type

A single event supports numeric or string value types only.

## Examples

The following are examples of single events:

- The time at which a particular batch file was processed.
- The checksum of an input file.
- The first and last CDR sequence number in a file (two separate statistic instances required).
- The list of billing output files that were generated from a collected CDR file (one statistic instance per billing output file).

## Groups

Groups are used to maintain a grouping of associated count type statistics whose counts can be checkpointed, committed and rolled back together. They are intended for use in batch environments where it is imperative that counts remain consistent across processing interruptions and resumptions mid-batch.

## Statistic Definitions

Statistic definitions are accessed via the Statistics Management node in the Management Center. There are four types of statistic definitions, namely:

- Custom Statistics—These are customised statistics which are highly specific to a customer's business logic. You can create, edit, and delete custom statistics.
- Product Statistics—These are the core statistics which are applicable to all deployments. Definitions for these statistics are created at installation time and they cannot be created or deleted. You can change the period (frequency at which they are collected), collection (if they are collected or not), persistence (if they are persisted or not), publication (if they are published or not), and lifetime (how long they are stored in the database) of the existing product statistic definitions.